

Sunet Drive



SUNET



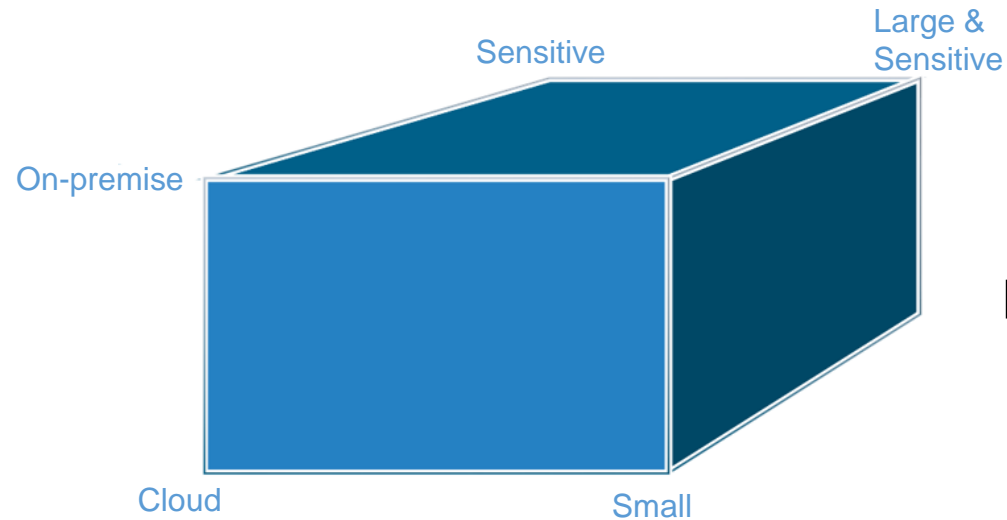
Status and plan for Swedens storage solution

- Research data – a model
- Data Integration
- Federated Data Storage and Governance
- High Level Architecture
- IdP Delegation
- Infrastructure as Code
- Monitoring and Operations
- Replication of data/backup
- Metadata integration



Research Data – A model

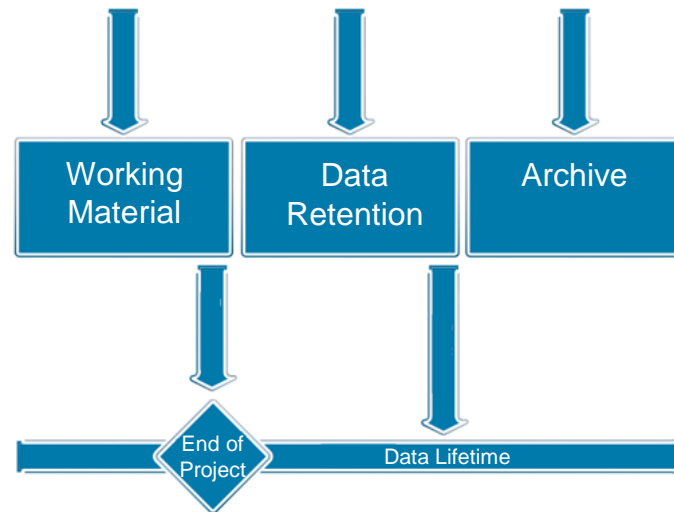
Classification



How do you access data?

Where is my data?

Lifecycle

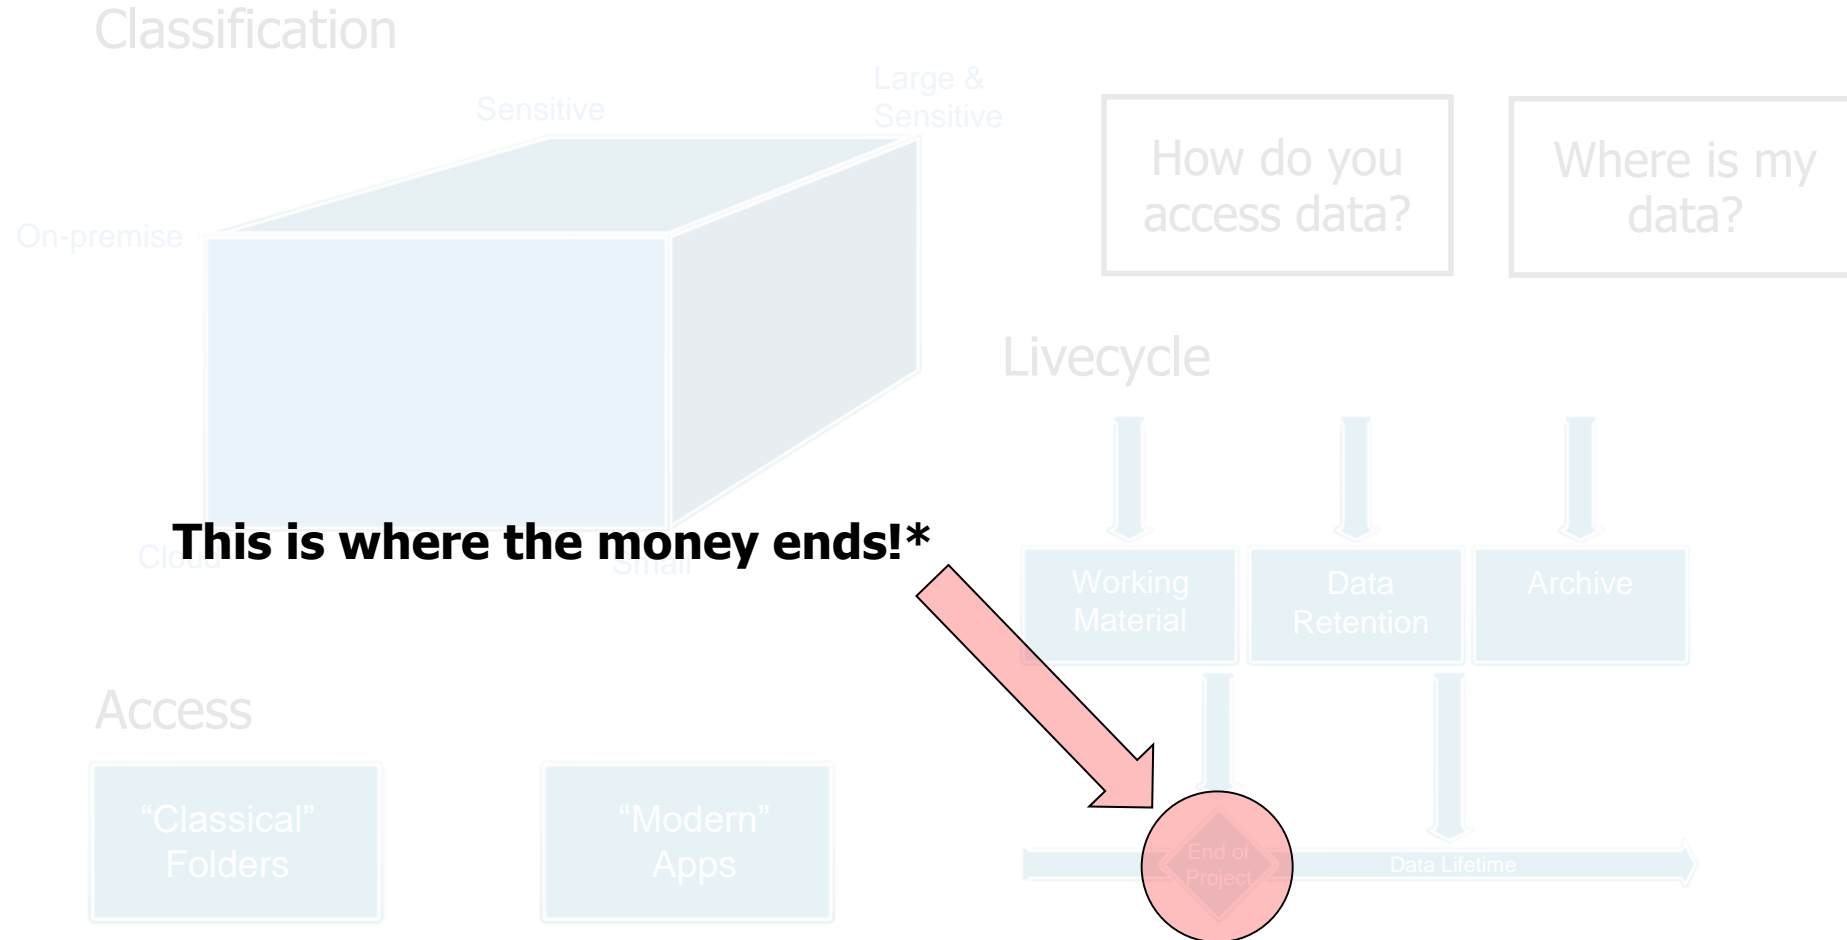


Access

“Classical”
Folders

“Modern”
Apps


Research Data – A model



*for many projects

Example - Integration

Primary data sources



Access

Access

- BoxSync - jusol-pukan@eduid.se@su.nextcloud.safedc.services
- GroupExternalStorage - user3@mx3.dev.nextcloud.safedc.services
- Nextcloud - richard@familybox.mfreitag.com
- Nextcloud - user3@mx3.dev.nextcloud.safedc.services

```
mount: can't find https://espace.simulo.its.uu.se/sites/UADM-DOKU/Documents in /etc/fstab
efss@EPSS-Linux:~$ sudo mount -t davfs https://espace.simulo.its.uu.se/sites/UADM-DOKU/Documents UADM-DOKU
Please enter the username to authenticate with server
https://espace.simulo.its.uu.se/sites/UADM-DOKU/Documents or hit enter for none.
Username: testforskada
Please enter the password to authenticate user testforskada with server
https://espace.simulo.its.uu.se/sites/UADM-DOKU/Documents or hit enter for none.
Password:
efss@EPSS-Linux:~$
```

Dropdepo

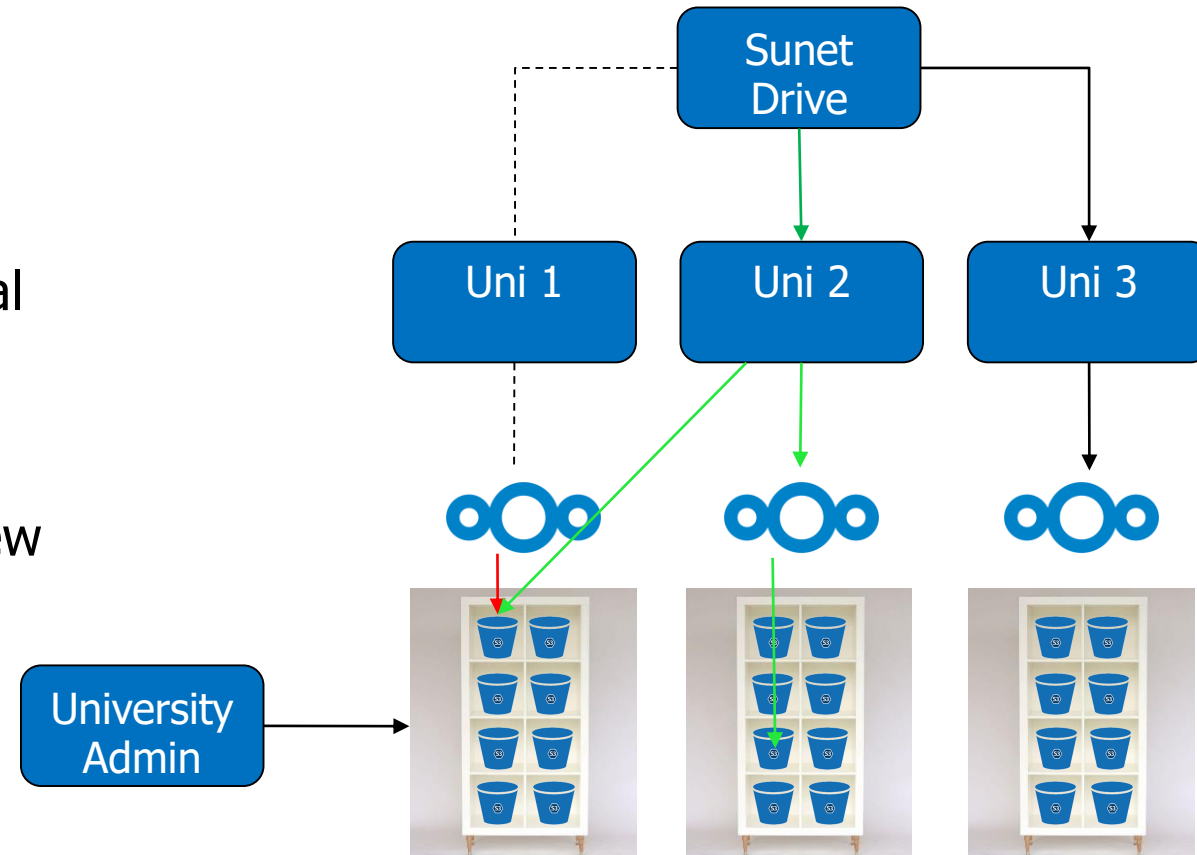
- 20180728_232921.mp4 14.6 MB, Jul 29
- fileupload.txt 19.8, Jul 30
- IMG_20180708_141743.jpg 1.1 MB, Jul 8
- nextcloud-10010002.apk 3.8 MB, Jul 29
- Screenshot_2_8-233410.jpg 302 KB, Jul 28
- Untitled.txt 3.8, Jul 17

18 files, 12 folders

The 'Access' section provides a visual overview of how the storage is accessed. It features a collage of images: a Windows File Explorer window showing a folder structure, a terminal window with mount and authentication commands, a smartphone displaying a file manager interface with a 'Nextcloud Manager' app, and a desktop application window showing a file list with details like file names, sizes, and dates. A list of Nextcloud accounts is also shown, including BoxSync, GroupExternalStorage, and Nextcloud users.

Federated Data Storage and Governance

- S3-buckets as main storage unit
- Nextcloud as access layer
- Individual nodes as part of global scale
- Co-management with Sunet
- Follow your data even from a new organisation
- **Separation of access vs ownership**



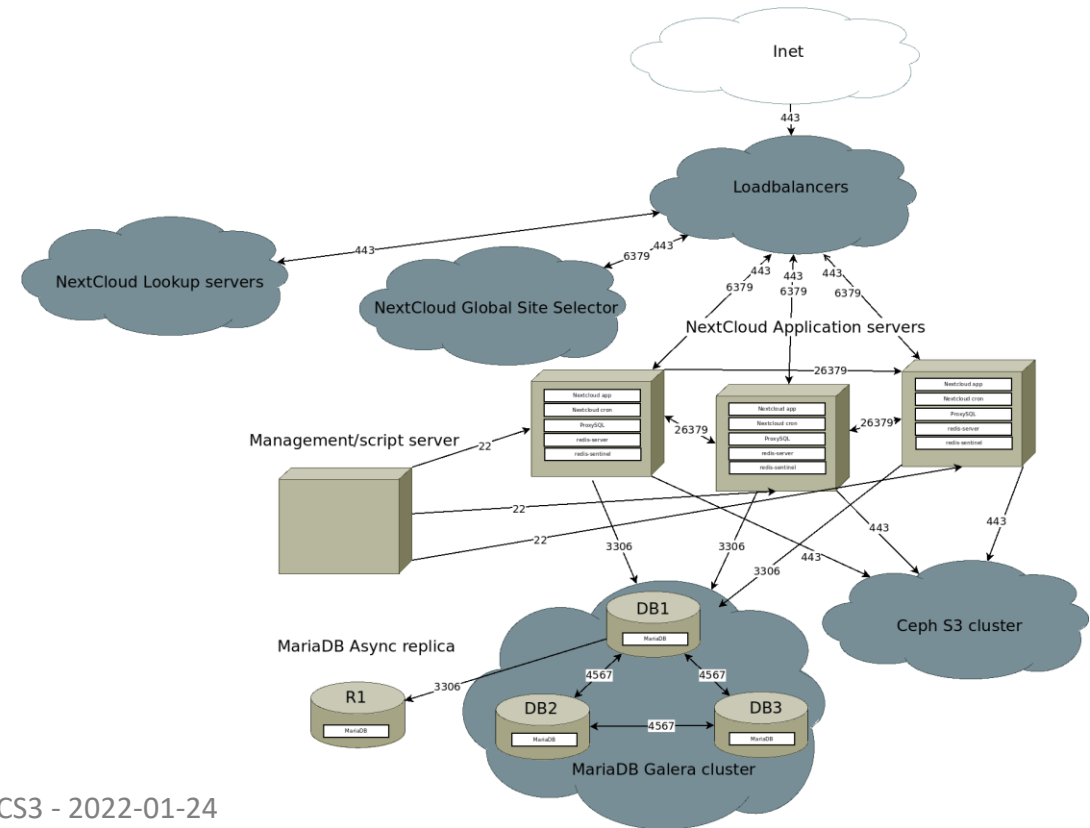
drive.sunet.se – High level architecture

Paying universities get an automatically provisioned HA node with no SPOF consisting of:

- Load-balanced Sunet Drive Frontend Servers
- MariaDB Galera cluster as database
- Redis Server/Sentinel Cluster for improved performance
- S3 storage integration
- Access to test environment
- Optional storage-mirroring
- Customer support

Non-Paying universities get:

- Access via a shared node
- Provisioning of user accounts
- ↓The option to upgrade to this↓ 😊






drive.sunet.se – IdP delegation

SWAMID/eduGAIN/Seamlessaccess



Choose Your Institution
Recent institutions

 SUNET sunet.se	>
 Uppsala University user.uu.se	>
 Stockholm University su.se	>

[Add another institution](#) [Edit](#)



sunet.drive.sunet.se



uu.drive.sunet.se



su.drive.sunet.se

EduID



eduID is easier and safer login.

EMAIL:
Richard@

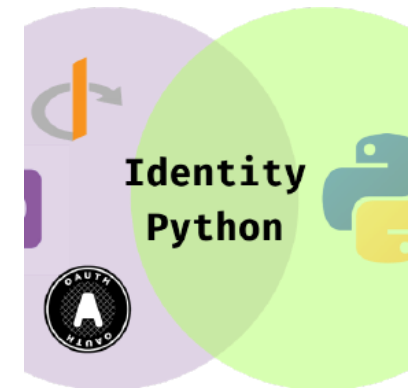
PASSWORD:
.....

[Forgot your password?](#)

SIGN IN



extern.drive.sunet.se



Infrastructure as Code*



Not Invented Here



Invented Here™



* <https://github.com/SUNET/multiverse/>

Shout Out to Thruk*

Nagios/Naemon/
cinga/Shinken
would not be the
same without
you!

* <http://www.thruk.org/>

Thruk

Current Network Status
Last Update: Fri Jan 21 10:05:26 UTC 2022 (≈90s)
Thruk 2.46.2
Logged in as nagiosadmin

Host Status Totals

Up	Down	Unreachable	Pending
292	0	0	0
All Problems		All Types	
0		292	

Service Status Totals

OK	Warning	Unknown	Critical	Pending
3576	0	0	0	0
All Problems		Unhand. Issues		All Types
0		0		3576

Service Status Details For All Host

1 2 3 4 5 6 7 8 9 10 ... 36 ▶ 100 per page

Select hosts / services with leftclick to send multiple commands. Select multiple with shift + mouse.

Host	Service	Status	Last Check	Duration	Attempt	Status Information
antagning.drive.sunet.se	Nextcloud status	OK	10:01:01	0d 0h 34m 25s	1/3	HTTP OK: HTTP/1.1 200 OK - 1769 bytes in 0.044 second resp...
	Nextcloud version	OK	10:01:01	0d 0h 34m 25s	1/3	HTTP OK: HTTP/1.1 200 OK - 1773 bytes in 0.058 second resp...
	PING	OK	10:01:45	16d 3h 3m 51s	1/4	PING OK - Packet loss = 0%, RTA = 1.36 ms
antagning.drive.test.sunet.se	Nextcloud status	OK	10:02:08	6d 17h 36m 48s	1/3	HTTP OK: HTTP/1.1 200 OK - 1778 bytes in 0.048 second resp...
	Nextcloud version	OK	10:02:31	6d 17h 36m 48s	1/3	HTTP OK: HTTP/1.1 200 OK - 1774 bytes in 0.036 second resp...
	PING	OK	10:02:54	45d 15h 18m 26s	1/4	PING OK - Packet loss = 0%, RTA = 0.81 ms
backup1.extern.drive.sunet.se	Active Users	OK	10:03:18	0d 23h 41m 34s	1/4	USERS OK - 0 users currently logged in
	Boot Disk	OK	10:03:41	87d 20h 21m 45s	1/4	DISK OK - free space: / 71505 MB (90% inode=98%):
	PING	OK	10:04:04	73d 2h 44m 8s	1/4	PING OK - Packet loss = 0%, RTA = 0.55 ms
	Packages available for upgrade	OK	10:04:27	0d 3h 33m 25s	1/4	APT OK: 0 packages available for upgrade (0 critical updates).
	Processes need restart	OK	10:04:50	7d 0h 27m 15s	1/4	OK - Services: none, Containers: none, Sessions: none
	Reboot Needed	OK	10:05:13	0d 23h 41m 12s	1/4	Reboot OK: No reboot required
	Root Disk	OK	10:03:05	43d 18h 3m 10s	1/4	DISK OK - free space: / 71505 MB (90% inode=98%):
	Scriptherder Status	OK	10:00:59	55d 20h 0m 19s	1/4	OK: 4 job in this state: cosmos[exit=0,age=58s], dockerhost_cleanup[exit=0,age...
	System Entropy	OK	10:01:22	85d 21h 25m 45s	1/4	OK: 1402 bytes in the pool.
	System Load	OK	10:01:45	87d 20h 19m 29s	1/4	OK - load average: 0.07, 0.05, 0.01
	System Memory	OK	10:02:08	87d 20h 19m 12s	1/4	MEMORY OK - 7144M free
	System NTP Time	OK	10:02:31	13d 20h 52m 48s	1/4	NTP OK: Offset -8.702278137e-06 secs
	Total Processes	OK	10:02:55	87d 20h 18m 41s	1/4	PROCS OK: 44 processes
Uptime	OK	10:03:18	0d 23h 41m 34s	1/4	OK: Linux backup1 5.4.0-96-generic - up 23 hours 40 minutes	
Var Disk	OK	10:03:41	73d 2h 44m 8s	1/4	DISK OK - free space: / 71505 MB (90% inode=98%):	

Replication of data/backup

Manual replication/backup strategy

- Based on rclone
- Manually mirror between data centers
- Backup-buckets with retention time
- Requires a lot of tuning

- Waiting for something better (deuxfleurs/garage?)



[[Website and documentation](#) | [Binary releases](#) | [Git repository](#) | [Matrix channel](#)]

Garage is a lightweight S3-compatible distributed object store, with the following goals:

- As self-contained as possible
- Easy to set up
- Highly resilient to network failures, network latency, disk failures, sysadmin failures
- Relatively simple
- Made for multi-datacenter deployments

Non-goals include:

- Extremely high performance
- Complete implementation of the S3 API
- Erasure coding (our replication model is simply to copy the data as is on several nodes, in different datacenters if possible)

Our main use case is to provide a distributed storage layer for small-scale self hosted services such as [Deuxfleurs](#).

SND/Doris Integration - Future



New data description

Title *

Enter a descriptive, preferably unique, name for the data description, both in Swedish and in English. If no title is available in Swedish, or if it is difficult to translate, the English title can be entered in both fields.

Data accessibility level *

Access to data through SND or other external actor *

Specify where data will be available. You choose whether the data is made available through SND's research data portal, via downloading or by request, or through an external s.. [Show more](#)

- Access to data through SND
- Access to data through an external actor

Level of accessibility *

Data that are shared by SND can have two different accessibility levels. Data should, in keeping with the recommendations from the Swedish Research Council, be openly accessible as .. [Show more](#)

- Data are freely accessible
- Access to data is restricted

Tack! Questions?



Micke Nordin

kano@sUNET.se



Richard Freitag

freitag@sUNET.se

